5 Sample Policy Areas

A series of questions will emerge from the need to support some combination of the Internet-based business requirements discussed above. What controls and procedures should be implemented to support our business needs? What type of risk profile do we fall under? What is our style and culture of doing business? Who is responsible? The elements that drive the answers to such questions form the policy framework for the organization.

The following sections contain hypothetical sample policy statements that address Internet-based security. The policy elements are derived from the major sources of security controls (e.g., software import control, encryption, and system architecture). The rationale that drives the selection of certain policy is given, followed by the actual sample policy statement(s), which are indented.

Each section contains multiple sample policies for use at the different risk profiles discussed in chapter 3. Some areas provide multiple examples at the same risk level to show the different presentation methods that might be used to get the message across.

Security policies fall into two broad categories: technical policies to be carried out by hardware or software, and administrative policy to be carried out by people using and managing the system. The following sections indicate each type of policy with an icon.

5.1 Identification and Authentication

Identification and Authentication (I&A) is the process of recognizing and verifying valid users or processes. I&A information is generally then used to determine what system resources a user or process will be allowed to access. The determination of who can access what should be part of a data categorization effort, described in section5.6****.

The chapter assumes that a decision has been made to allow connectivity to internal systems from the Internet. If there is no connectivity, there is no need for I&A. Many organizations separate Internet-accessible systems from internal systems through the use of firewalls and routers. See Sections *** on architecture and firewalls.

Authentication over the Internet presents several problems. It is relatively easy to capture identification and authentication data (or any data) and replay it in order to impersonate a user. As with other remote I&A, and often with internal I&A, there can be a high level of user dissatisfaction and uncertainty which can make I&A data obtainable via social engineering. Having additional I&A for use of the Internet may also contribute to I&A data proliferation which is difficult for users to manage. Another problem is the ability to hijack a user session after the I&A has been performed.

There are three major types of authentication available: static, robust, and continuous. Static authentication includes passwords and other techniques that can be compromised through replay attacks. They are often called reusable passwords. Robust authentication involves the use of cryptography or other techniques to create one-time passwords that are used to create sessions. These can be compromised by session hijacking. Continuous authentication prevents session hijacking.

1. Static Authentication

Static authentication only provides protection against attacks in which an imposter cannot see, insert or alter the information passed between the claimant and the verifier during an authentication exchange and subsequent session. In these cases, an imposter can only attempt to assume a claimant's identity by initiating an access control session as any valid user might do and trying to guess a legitimate user's authentication data. Traditional password schemes provide this level of protection, and the strength of the authentication process is highly dependent on the difficulty of guessing password values and how well they are protected.

2. Robust Authentication

This class of authentication mechanisms relies on dynamic authentication data that changes with each authenticated session between a claimant and verifier. An imposter who can see information passed between the claimant and verifier may attempt to record this information, initiate a separate access control session with the verifier, and replay the recorded authentication data in an attempt to assume the claimant's identity. Level 1 strong authentication protects against such attacks, because authentication data recorded during a previous session will not be valid for any subsequent sessions.

However, robust authentication does not provide protection against active attacks in which the imposter is able to alter the content or flow of information between the claimant and verifier after they have established a legitimate session. Since the verifier binds the claimant's identity to the logical communications channel for the duration of the session, the verifier believes that the claimant is the source of all data received through this channel.

Traditional fixed passwords would fail to provide robust authentication because the password of a valid user could be viewed and used to assume that user's identity later. However, one-time passwords and digital signatures can provide this level of protection.

3. Continuous Authentication

This type of authentication provides protection against impostors who can see, alter, and insert information passed between the claimant and verifier even after the claimant/verifier authentication is complete. These are typically referred to as active attacks, since they assume that the imposter can actively influence the connection between claimant and verifier. One way to provide this form of authentication is to apply a digital signature algorithm to every bit of data that is sent from the claimant to the verifier. There are other combinations of cryptography that can provide this form of authentication but current strategies rely on applying some type of cryptography to every bit of data sent. Otherwise, any unprotected bit would be suspect.

5.1.1 General Internet I&A Policies

Although passwords are easily compromised, an organization may find that a threat is not likely, would be fairly easy to recover from, or would not affect critical systems (which may have separate protection mechanisms).

Low

Authentication is required for access to corporate systems from the Internet. The minimum standard for authentication is passwords as described in ****.

Medium

Internet access to XYZ category of information and processing (low impact if modified, unavailable, or disclosed) requires a password and access to all other resources requires robust authentication.

Telnet access to corporate resources from the Internet requires the use of robust authentication.

High

Internet access to all systems behind the firewall requires robust authentication. Access to ABC information and processing (high impact if modified, unavailable, or disclosed) requires continuous authentication.

5.1.2 Password Management Policies

The following are general password policies applicable for Internet use:

Passwords and user logon IDs will be unique to each authorized user.

Passwords will consist of a minimum of 6 alphanumeric characters (no common names or phrases). There should be computer-controlled lists of proscribed password rules and periodic testing (e.g., letter and number sequences, character repetition, initials, common words, and standard names) to identify any password weaknesses.

Passwords will be kept private i.e., not shared, coded into programs, or written down

Passwords will be changed every 90 days (or other period). Most systems can enforce password change with an automatic expiration and prevent repeated or reused passwords.

User accounts will be frozen after 3 failed logon attempts. All erroneous password entries will be recorded in an audit log for later inspection and action, as necessary.

Sessions will be suspended after 15 minutes (or other specified period) of inactivity and require the password to be reentered.

Successful logons should display the date and time of the last logon and logoff.

Logon IDs and passwords should be suspended after a specified period of disuse.

For high risk systems:

After excessive violations, the system should generate an alarm and be able to simulate a continuing session (with dummy data, etc.) for the failed user (to keep this user connected while personnel attempt to investigate the incoming connection.

5.1.3 Robust Authentication Policy

The decision to use robust authentication requires an understanding of the security gained and the cost of user acceptance and administration. User acceptance will be dramatically improved if users are appropriately trained why the robust authentication is being used and how to use it.

There are many technologies available that provide robust authentication including dynamic password generators, cryptography-based challenge/response tokens and software, and digital signatures and certificates. If digital signatures and certificates are used another policy area is opened up – what are the security requirements for the certificates?

Users of robust authentication must receive training prior to use of the authentication method.

> *Employees are responsible for safe handling and storage of all company authentication devices. Authentication tokens should not be stored with a computer that will be used to access corporate systems. If an authentication device is lost or stolen, the loss must be immediately reported to security so that the device can be disabled.*

5.1.4 Digital Signatures and Certificates

If I&A makes use of digital signatures, then certificates are required. They can be issued by the organization or by a trusted 3rd party. Commercial PKI infrastructures are emerging within the Internet community. Users can obtain certificates with various levels of assurance.

Example of Different Levels for a PKI

- Level 1 certificates verify electronic mail addresses. This is done through the use of a personal information number that a user would supply when asked to register. This level of certificate may also provide a name as well as an electronic mail address; however, it may or may not be a genuine name (i.e., it could be an alias).
- Level 2 certificates verify a user's name, address, social security number, and other information against a credit bureau database.
- Level 3 certificates are available to companies. This level of certificate provides photo identification (e.g., for their employees) to accompany the other items of information provided by a Level 2 certificate.

Once obtained, digital certificate information may be loaded into an electronic mail application or a web browser application to be activated and provided whenever a web site or another user requests it for the purposes of verifying the identity of the person with whom they are communicating. Trusted certificate authorities are required to use such systems effectively, otherwise fraudulent certificates could easily be issued.

Many of the latest web servers and web browsers incorporate the use of digital certificates. Secure Socket Layer (SSL) is the technology used in most Web-based applications. SSL version 2.0 supports strong authentication of the Web server, while SSL 3.0 added client side authentication. Once both sides are authenticated, the session is encrypted, provide protection against both eavesdropping and session hijacking. The

digital certificates used are based on the X.509 standard and describe who issued the certificate, the validity period, and other information.

Oddly enough, passwords still play an important role even when using digital certificates. Since a digital certificate is stored on a computer, they can only be used to authenticate the computer, rather than the user, unless the user provides some other form of authentication to the computer. Passwords or passphrases are generally used; smart cards and other hardware tokens will be used in the future.

> *Any company systems making limited distribution data available over the Internet shall use digital certificates to validate the identity of both the user and the server. Certificates may only be issued by COMPANY-approved Certificate Authorities. Certificates at the user end will be used in conjunction with standard technologies such as Secure Sockets Layer to provide continuous authentication to eliminate the risk of session hijacking.*
>
> *Access to digital certificates stored on personal computers should be protected by passwords or passphrases. All policies for password management must be followed.*

5.2 Software Import Control

Data on computers is rarely static. Mail arrives and is read. New applications are loaded from floppy, CDROM, or across a network. Web-based interactive software downloads executables that run on a computer. Each modification runs the risk of introducing viruses, damaging the configuration of the computer, or violating software-licensing agreements. Organizations need to protect themselves with different levels of mechanisms depending on the sensitivity to these risks.

Software Import Control provides an organization with several different security challenges:

- Virus and Trojan Horse Prevention, Detection, and Removal
- Controlling Interactive Software (Java, ActiveX)
- Software Licensing

Each challenge can be categorized according to the following criteria:

- Control - who initiates the activity, and how easily can it be determined that software has been imported
- Threat type - executable program, macro, applet, violation of licensing agreement
- Cleansing Action - scanning, refusal of service, control of permissions, auditing, deletion

When importing software onto a computer one runs the risk of getting additional or different functionality than one bargained for. The importation may occur as a direct action, or as a hidden side-effect which is not readily visible. Examples of direct action are:

- File Transfer - utilizing FTP to transfer a file to a computer
- Reading E-mail - causing a message which has been transferred to a computer to be read, or using a tool (e.g., Word) to read an attachment

- Downloading software, from a floppy disk or over the network can spawn indirect action. Some examples include:
- Reading a Web page which downloads a Java applet to your computer
- Executing an application such as Microsoft Word, and opening a file infected with a Word Macro Virus.

Viruses imported on floppy disks or infected vendor media will continue to be a major threat.

5.2.1 Virus Prevention, Detection, and Removal

The most common “carrier” of viruses has been the floppy disk, since “sneaker net” was the most common means of transferring software between computers. As telephone-based bulletin boards became popular, viruses traveled more frequently via modem. The Internet provides yet another channel for virus infections, one that can often bypass traditional virus controls.

For organizations that allow downloading of software over the Internet (which can be via Internet email attachments) virus scanning at the firewall can be an appropriate choice - but it does not eliminate the need for client and server based virus scanning, as well. For several years to come, viruses imported on floppy disks or infected vendor media will continue to be a major threat.

A virus is a self-replicating program spread from executables, boot records, and macros. Executable viruses modify a program to do something other than the original intent. After replicating itself into other programs, the virus may do little more than print an annoying message, or as damaging as deleting all of the data on a disk. There are different levels of sophistication in how hard a virus may be to detect.

Simple viruses can be easily recognized by scanning for a signature of byte strings near the entry point of a program, once the virus has been identified. Polymorphic viruses modify themselves as they propagate, therefore have no signature, and can only be found (safely) by executing the program in a virtual processor environment. Boot record viruses modify the boot record such that the virus is executed when the system is booted.

Applications that support macros are at risk for macro viruses. Macro viruses are commands that are embedded in data. Vendor applications, such as Word, Excel, or printing standards such as Postscript are common targets. When the application opens the data file the infected macro virus is instantiated.

The security service policy for viruses has three aspects:

- Prevention - policies which prevent the introduction of viruses into a computing environment
- Detection - determination that an executable, boot record, or data file is contaminated with a virus
- Removal - deletion of the virus from the infected computing system may require reinstallation of the OS from the ground up, deleting files, or deleting the virus from an infected file.

There are various factors that are important in determining the level of security concern for virus infection of a computer. Viruses are most prevalent on DOS, Windows (3.x, 95), and NT operating systems. However there are also some UNIX and even LINUX viruses.

The frequency that new applications or files are loaded on to the computer is proportional to the susceptibility of that computer to viruses. Configuration changes resulting from exposure to the Internet, exposure to mail, or receipt of files from external sources are more at risk for contamination.

The greater the value of the computer or data on the computer, the greater the concern should be for insuring that virus policy as well as implementation procedures are in place. The cost of removal of the virus from the computing environment must be considered within your organization as well as from customers you may have infected. Cost may not always be identified as monetary; company reputation and other considerations are just as important.

It is important to note that viruses are normally introduced into a system by a voluntary act of a user (e.g., installation of an application, FTP of a file, reading mail, etc.) Prevention policies can therefore focus on limiting introduction of potentially infected software and files to a system. This also indicates in a high-risk environment that virus-scanning efforts should be focused on when new software or files are introduced to maximize protection.

Low Risk

Software import control policies for low risk environments should concentrate on educating users on their responsibilities for regularly scanning for viruses.

Prevention:

> *Users will be trained about the possibility of receiving viruses and other malicious code from the Internet and on the use of virus scanning tools.*

Detection:

> *Off the shelf virus-scanning tools will be used to scan computers weekly. No auditing of virus scanning tool records is necessary.*
>
> *Employees will inform the system administrator of any virus detected, configuration change, or different behavior of a computer or application. When informed that a virus has been detected the system administrators will inform all users with access to the same programs or data that a virus may have also infected their system. The users will be informed of the steps necessary to determine if their system is infected and the steps to take to remove the virus. Users will report the results of scanning and removal activity to the system administrators.*

Removal:

> *Any machine thought to be infected by a virus will immediately be disconnected from all networks. The machine will not be reconnected to the network until system administration staff can verify that the virus has been removed. When applicable, off-*

the-shelf virus scanning tools will be used to remove a virus from an infected file or program. If virus-scanning software fails to remove the virus, all software on the computer will be deleted including boot records if necessary. The software will then be reinstalled uninfected sources and re-scanned for viruses.

Medium Risk

Software import control policies for medium risk environments should dictate more frequent scanning for viruses, and the use of server and email virus scanners.

Prevention:

Software will be downloaded and installed only by network administrators (who will scan or test software).

Anti-virus software will be installed in file servers to limit the spread of viruses within the network. Scanning of all files and executables will occur daily on the file servers. Workstations will have memory resident anti-virus software installed and configured to scan data as it enters the computer. All incoming electronic mail will be scanned for viruses. Programs will not be executed, and files opened by applications prone to macro viruses without prior scanning.

Employee security training will include information about virus infection risks:

Virus scanning software is limited to the detection of viruses that have been previously identified. New and more sophisticated viruses are being developed constantly. Virus scanning software must be updated on a regular (monthly or quarterly) basis to maintain currency with the latest viruses. It is important to inform the system administrator of any different or out of the ordinary behavior a computer or application exhibits. It is important to immediately disconnect a computer that is infected or thought to be infected from networks to reduce the risk of spreading a virus.

Detection:

Off the shelf virus-scanning tools will be used to scan computers on a daily basis. The virus scanning tools will be updated on a monthly basis. All software or data imported onto a computer (from floppy disk, e-mail, or file transfer) will be scanned before being used.

Virus scanning logs will be recorded, reported and examined by the system administration staff. Employees will inform the system administrator of any virus that is detected, configuration change or different behavior of a computer or applications.

When informed that a virus has been detected the system administrators will inform all users who may have access to the same programs or data that a virus may have also infected their system. The users will be informed of the steps necessary to determine if their system is infected and the steps to take to remove the virus. Users will report the results of scanning and removal activity to the system administrators.

Removal:

Any machine thought to be infected by a virus will immediately be disconnected from all networks. The machine will not be reconnected to the network until system administration staff can verify that the virus has been removed. When applicable, off-

the-shelf virus scanning tools will be used to remove a virus from an infected file or program. If virus-scanning software fails to remove the virus, all software on the computer will be deleted including boot records if necessary. The software will then be reinstalled uninfected sources and rescanned for viruses.

High Risk

High security level systems contain data and applications that are critical to the corporation. Infection will cause considerable loss of time, data, and potentially harm the reputation of the corporation. Large numbers of computers may be involved. All reasonable measures possible to prevent virus infection must be taken.

Prevention:

The CIO/Security Administrator must approve all applications before they can be installed on a computer. No unauthorized applications may be installed on a computer. Software configurations will be scanned on a monthly basis to validate that no extraneous software has been added to a computer.

Software will be installed only from approved internal servers to limit exposure to contaminated software. No software will be downloaded from the Internet onto any computer. File transfer "gets" from external sources will not be permitted.

Anti-virus software will be installed in file servers to limit the spread of viruses within the network. Scanning of all files and executables will occur daily on the file servers. Workstations will have memory resident anti-virus software installed and configured to scan data as it enters the computer. Programs will not be executed, nor files opened by applications prone to macro viruses without prior scanning.

All incoming mail and files received from across a network must be scanned for viruses as they are received. Virus checking will be performed if applicable at firewalls that control access to networks. This will allow centralized virus scanning for the entire organization, and reduce overhead by simultaneously scanning incoming messages that have multiple destinations. It also allows for centralized administration of the virus scanning software, limiting the locations on which the latest virus scanning software needs to be maintained.

Employee security training will include the following information about virus infection risks:

Virus scanning software is limited to the detection of viruses that have been previously identified. New viruses and more sophisticated viruses are being developed constantly. Virus scanning software must be updated on a regular (monthly or quarterly) basis to maintain currency with the latest viruses. It is important to inform the system administrator of any different or out of the ordinary behavior a computer or application exhibits. It is important to immediately disconnect a computer that is infected or thought to be infected from networks o reduce the risk of spreading a virus.

Failure to follow these policies may result in punishment according to company standards.

Detection:

All software must be installed on a testbed and tested for viruses before being allowed on an operational machine. Only after receiving approval from the CIO/Security Administrator may software be moved to operational machines.

Use off-the-shelf scanning software will be enhanced by state of the art virtual machine emulation for polymorphic virus detection. All other new virus detection methods will be incorporated into the detection test bed. To keep abreast of the latest viruses which have been identified, scanning software will be updated monthly or as updates arrive.

Virus scanning of all file systems on a daily basis is mandatory. Virus scanning results will be logged, automatically collected, and audited by the system administration staff.

All data imported on a computer (from floppy disk, e-mail, or file transfer) will be scanned before being used. Employees will inform the system administrator of any virus that is detected, configuration change, or different behavior of a computer or application.

When informed that a virus has been detected the system administrators will inform all users who may have access to the same programs or data that a virus may have also infected their system. The users will be informed of the steps necessary to determine if their system is infected and the steps to take to remove the virus.

Removal:

Any machine thought to be infected by a virus will immediately be disconnected from all networks. The machine will not be reconnected to the network until system administration staff can verify that the virus has been removed. When applicable, off-the-shelf virus scanning tools will be used to remove a virus from an infected file or program. If virus-scanning software fails to remove the virus, all software on the computer will be deleted including boot records if necessary. The software will then be reinstalled uninfected sources and rescanned for viruses.

5.2.2 Controlling Interactive Software

A programming environment evolving as a result of Internet technology is Interactive Software, as exemplified by Java and ActiveX. In an Interactive Software environment a user accesses a server across a network. The server downloads an application (applet) onto the user's computer that is then executed. There are significant risks involved in this strategy. Fundamentally, one must trust that what is downloaded will do what has been promised.

There have been various claims, now effectively debunked, that when utilizing languages such as Java it is impossible to introduce a virus because of restrictions within the scripting language for file system access and process control. Because of these significant risks, there are several levels of trust that a user must have before employing this technology:

- The server can be trusted to download trustworthy applets;
- The applet will execute in a limited environment restricting disk reads and writes to functions which do not have security;
- The applet can be scanned to determine if it is safe;

- Scripts are interpreted, not precompiled. As such, there is a risk that a script can be modified in transit and not perform its original actions.

Java and ActiveX Security Models

Java is a programming language developed by Sun Microsystems in order to provide a mechanism for allowing software programs to be downloaded over the Internet and run on a variety of workstations and personal computers. Java is interpreted at run time, and actually runs on top of software called the Java Virtual Machine. The JVM runs on Unix, Windows, Macintosh, or other operating systems to allow Java applets to run identically across heterogeneous environments.

Java's security model is to tightly control the environment in which applets can operate, creating a safe "sandbox" for applet processing. Applets can only communicate with the server from which they were downloaded and are prohibited from accessing local disks or network connections. However, many bugs have been discovered in Java that allow clever programmers to create applets that can easily escape the sandbox. Sun has responded by making the walls of the sandbox higher, but new vulnerabilities are still regularly found.

ActiveX is an outgrowth of Microsoft's Object Linking and Embedding technology, which allows programs to communicate with functions within standard applications, such as word processors and spreadsheets. ActiveX allows applications to communicate over the Internet, and for applets to be downloaded to a user machine and access local resources.

Microsoft's ActiveX security model is quite different than Sun's Java model. ActiveX allows the user to specify the trust relationship with the server that downloads the applet. If the server is trusted, and its identity verified by the use of digital certificates, an ActiveX applet can be downloaded and operate much like any other piece of software on the user computer. Digital signatures can be used, called Authenticode, which verifies that the code came from a trusted developer and was not modified.

Neither security model is clearly superior. The Java approach limits the damage a malicious applet can cause - if all bugs in Java are found and addressed. The ActiveX approach mirrors the way businesses buy and install commercial software, but it places a large degree of authority with the individual user. Some firewall vendors support blocking or authenticating applets at the firewall.

It should be noted that a user may not be aware that an applet is being downloaded and executed on his/her computer. As such, security measures must be set up in advance to prevent such an occurrence.

Low Risk

Users should be trained about the risks of Interactive Software and how to configure their browsers to prevent downloading of applets.

Medium Risk

If possible, firewalls will be configured to block the reception of applets from external sources and block the distribution of applets outside of internal networks unless authentication technology is used to protect it from untrusted sources.

Users will configure their browsers to accept applets only from trusted servers. If this not possible, then browsers will be configured to not accept applets.

If appropriate, use of applets will be restricted only to development networks and not permitted on operational networks.

High Risk

Use of Interactive Software is prohibited. Web browsers, and where applicable, firewalls, will be configured and the configurations audited by the system administration staff to inhibit the downloading of applets. Failure to comply with this policy will result in disciplinary action against the employee.

5.2.3 Software Licensing

The Internet has allowed many software companies to use new means of distributing software. Many companies allow the downloading of trial versions of their products, sometimes limited versions ("crippleware") or versions that only operate for a limited period of time. However, many companies take a shareware approach, allowing fully functional copies of software to be downloaded for trial use and requiring the user to register and pay for the software when used for commercial purposes.

When users forget or decline to properly register software downloaded over the Internet, the company can be in violation of software licenses. This may put a company at severe risk of penalties, or loss of reputation if discovered. The Business Software Alliance and the Software Publishing Alliance actively audit corporate licensing and pursue violators. Internet security policy should detail corporate policy on downloading commercial software.

Low Risk

Vendor licensing regulations will be followed for all commercial software downloaded over the Internet. Trial versions of programs should be deleted after the trial period, or the software should be procured through approved procedures.

Medium/High Risk

Commercial software should not be downloaded over the Internet without approval of system administration. All software to be used on COMPANY computers can only be installed by system administration, following all licensing agreements and

procedures. The system administration staff will inspect the computers periodically to verify that only approved and licensed software has been installed. Violation of this policy may result in disciplinary action.

5.3 Encryption

Encryption is the primary means for providing confidentiality services for information sent over the Internet. Encryption can be used to protect any electronic traffic, such as mail messages or the contents of a file being downloaded. Encryption also can protect information in storage, such as in databases or stored on computer systems where physical security is difficult or impossible (such as on laptop computers that may be left in hotel rooms).

There are a number of management issues related to encryption use:

The United States government currently imposes export controls on strong cryptography, currently defined as any encryption system with an encryption key longer than 40 bits that does not provide for key recovery. There are no restrictions on encryption for domestic use. For use in foreign countries, or in networks that include nodes in a foreign country, each use must receive export approval. In addition, some countries such as France and China impose their own set of controls on domestic use of encryption. Any use of encryption involving foreign countries needs to be thoroughly researched.

Management ability to monitor internal communications or audit internal computer systems may be impacted by the use of encryption. If employees encrypt outgoing email messages, or the content of the hard drive on their desktop computer, system administrative personnel will be unable to audit such messages and files. In addition, if the decryption keys are lost or damaged, the data may be permanently lost. If this level of monitoring or disaster recovery is important to your business operations, policy should mandate the use of encryption systems that support some form of key recovery.

There are a bewildering array of encryption algorithms and encryption key lengths. Policy should mandate the use of algorithms that have been in use commercially long enough to provide some assurance of security. The length of the encryption keys used should be driven by the value of the data to be encrypted. In general, given the state of the technology, encryption using keys of 40 or fewer bits is only acceptable for use behind the firewall - for keeping the honest people honest. Leading cryptographers recommend businesses use key lengths of at least 75 bits, with 90 bits being preferable. The Data Encryption Standard uses 56 bit keys, which is still acceptable for near term use. NIST is developing a new standard, the Advanced Encryption Standard. Triple DES, which provides an effective key length of 112 bits, is recommended in the interim.

The NIST Computer Security Handbook provides more detail on encryption and cryptography.

5.3.1 General Encryption Policy

To assure interoperability and consistency across the enterprise, corporate policy should mandate standards to which encryption systems must comply, specifying algorithms and parameters to be used. To assure interoperability and reduce life cycle costs, standard products should be selected for corporate use. While many easy to use and strong

COTS products are available, organizations still need to understand the overhead involved with encryption. The secure generation, storage, and transmission of keys as well as ensuring interoperability and, if needed, key recovery will require significant resources. (In general, key recovery is not an issue for Internet use, except for very high-risk organizations, but the same encryption will probably be used for local storage.)

Medium - high

Encryption should be used for (list specific types of information) that will be stored in non-secure locations or transmitted over open networks such as the Internet. Any encryption of other company information must be approved in writing. Where encryption is used, company-approved standard algorithms and standard products must be used. The minimal encryption key length for sensitive or confidential data is 56 bits - 75 bit keys are recommended.

The security of any encryption system is very dependent on the secrecy of the encryption keys used - procedures for secure generation and management of those keys are critical.

(Specific office) is responsible for developing procedures for encryption use and for training users.

Low - Medium

Encryption keys should be considered synonymous with the company's most sensitive category of information and access to those keys must be restricted on a need to know basis. The keys to be used for encryption must be generated by means that are not easily reproducible by outside parties.

Medium-High

The use of hardware-based random number generators is required for generating keys for encryption of company information. Use of a software-based random number generator will require written approval. Encryption keys should be considered sensitive information and access to those keys must be restricted on a need to know basis.

For symmetrical or secret key systems, such as DES, keys must only be transmitted via secure means. Since any compromise of these secret keys will make any use of encryption useless, security policy must detail acceptable means of distributing such keys.

When encryption is to be used, secure means must be used for all distribution of secret keys. Acceptable approaches include:

- *Use of public key exchange algorithms*
- *Double wrapped internal mail*
- *Double wrapped courier mail*

Encryption keys must not be sent via electronic mail, unless the electronic mail is encrypted using keys that have been previously exchanged using secure means. The keys used to encrypt information must be changed at the same frequency as the passwords used to access the information.

Data that has been encrypted can be lost forever if the decryption key is lost or damaged. Given that encryption will generally be used to protect valuable information, the loss of decryption keys can result in considerable damage. A number of approaches to assure accessibility of decryption keys have been proposed and are appearing in commercial products. As this technology is in its early stage, there are likely to be interoperability issues between products - corporate security staff should maintain a list of acceptable technologies or products.

All encryption products used must support some form of technology to make encryption keys available to management for all encryption of stored company information. Where encryption is used, company-approved key recovery implementations must be used. Any use of encryption without such technology must be approved in writing.

The use of public key encryption technologies requires both a public key and a private key to be created and maintained for each user. Public keys must be distributed or stored in a manner that they are accessible to all users. In advanced applications, digital certificates may be used to distribute public keys via Certificate Authorities. Secret keys are similar to passwords, and must be kept private by each user. Organizations may choose to require that all employees' secret keys be available to management.

Low

COMPANY will maintain listings of public keys for authorized encryption users. These lists will be stored on the authentication server or distributed via email. User secret keys must be kept private and should be treated the same as passwords.

Medium-High

The COMPANY Certificate Authority server will maintain all valid public keys for authorized encryption users. For secure communications with external entities, the COMPANY will accept Digital Certificates only from approved Certificate Authorities.

User secret keys must be kept private and should be treated the same as passwords. Any potential compromise of a user secret key must be reported to the security department immediately.

5.3.2 Remote Access

Organizations use both dial-in lines and the Internet to provide remote access. Both types of connections can be monitored but it is much more likely for Internet connections. The entity monitoring a session can read all the traffic including downloaded or uploaded files, email, and authentication information. The problem of monitoring accounts to gain authentication is discussed in the Identification and Authentication chapter. If organizations encrypt the entire session, it will address both the authentication and confidentiality issues. For medium-high security environments, encryption can be used to prevent unauthorized viewing of the information flowing during remote access connections.

Medium-high

All remote access to COMPANY computer systems, whether via dial-up or Internet access, must use encryption services to protect the confidentiality of the session.

COMPANY approved remote access products must be used to assure interoperability for remote access server encryption technologies.

Information regarding access to company computer and communication systems, such as dial-up modem phone numbers, is considered confidential. This information must not be posted on electronic bulletin boards, listed in telephone directories, placed on business cards, or made available to third parties without the written permission of the Network Services Manager. The Network Services Manager will periodically scan direct dial-in lines to monitor compliance with policies and may periodically change the telephone numbers to make it more difficult for unauthorized parties to locate company communications numbers.

5.3.3 Virtual Private Networks

Encryption is also used to create virtual private networks on the Internet. This is discussed in the Architecture chapter.

5.4 System/Architecture Level

Connecting to the Internet will involve a number of system architectural decisions that impact overall system security. One initial consideration will be the architecture of the firewall and the connection to the Internet - this is discussed in section 5.2.5. Other architectural choices that require policy decisions include the use of the Internet to connect physically separate networks (Virtual Private Networks,) providing remote access to systems connected to the Internet, and provided access to internal databases from the Internet. The security issues and sample policy statements related to these areas are detailed below.

5.4.1 Virtual Private Networks

Many organizations have local area networks and information servers spread across multiple locations. When organization-wide access to information or other LAN-based resources is required, leased lines are often used to connect the LANs into a Wide Area Network. Leased lines are relatively expensive to set up and maintain, making the Internet an attractive alternative for connecting physically separate LANs.

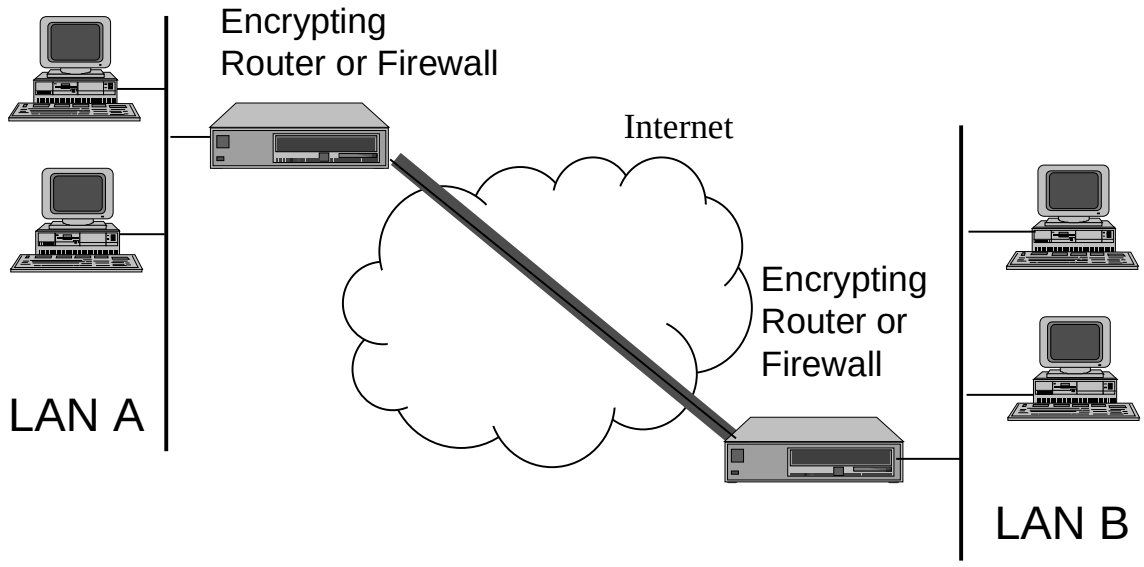

Figure 5-1 Virtual Private Networks can be established in a variety of configurations:

The major shortcoming to using the Internet for this purpose is the lack of confidentiality of the data flowing over the Internet between the LANs, as well as the vulnerability to spoofing and other attacks. Virtual Private Networks use encryption, as discussed in section 5.4.1, to provide the required security services. Typically encryption is performed between firewalls, and secure connectivity is limited to a small number of sites.

Security isn't the only issue with using the Internet to connect LANs. The Internet today provides no performance or reliability guarantee. Files or messages may be delayed or not delivered at all, depending on the overall state of the network and the state of individual routers, servers, and networks that make up the Internet.

Medium-high

Virtual Private Networks between sites must not use the Internet to carry time critical traffic. Where the level of reliability typically provided by the Internet is not sufficient to guarantee the required level of service to users, other means of interconnection must be used.

High

When the Internet is used to provide Virtual Private Network connections between sites, means of rapidly providing backup connections must be maintained to return service in the event of an Internet outage or denial of service.

One important consideration when creating Virtual Private Networks is that the security policies in use at each site must be equivalent. A VPN essentially creates one large network out of what were previously multiple independent networks. The security of the VPN will essentially fall to that of the lowest common denominator - if one LAN allows unprotected dial-up access, all resources on the VPN are potentially at risk.

The establishment of Virtual Private Networks (VPNs) over the Internet between company networks requires written approval of the CIO. Adding networks to an existing VPN also requires written approval of the CIO. A review and update of the security policies in use at each site to be connected to the VPN must be performed before operation will be authorized.

Trusted Links - The firewall encrypts all traffic destined for the remote host or network and decrypts all traffic it receives from. Traffic flows between hosts in a Trusted VPN relationship freely, as if there were no firewalls in between. The traffic is effectively routed by the firewalls involved, bypassing the proxies and thus not requiring any authentication at the firewall itself. Any two hosts who are part of a VPN Trusted Link have full network connectivity between them, and may communicate using any TCP/IP services that they support. Trusted Links are often used to connect geographically separate networks belonging to the same organization, each with their own connections to a local Internet service provider, into a seamless single virtual network in a secure fashion.

Private Links - The traffic is encrypted between the firewall and the remote host or network just as it is for the Trusted Link. However, traffic from remote hosts in a Private Link relationship is not freely routed, but must be proxied by the firewall and connections authenticated there as dictated by the firewall's usual proxy access policies. This relationship provides authentication of the network source of the traffic and confidentiality for the data, but the two networks maintain distinct network security perimeters, and only services which the firewall is configured to proxy can be used through it. Private Links are often used between the networks of organizations that do not wish to allow full access between their networks, but need confidentiality of the traffic between them.

Pass-through Links - Pass-through links are used to forward encrypted traffic between hosts on opposite sides of the firewall who are members of their own VPN peer relationship. This allows a firewall situated between two other VPN peers to be configured to route that encrypted data across. The intermediate firewall does not decrypt this traffic, nor does it need to know the encryption key used, it merely needs to know the addresses of the hosts on both sides of the link so it knows to allow the encrypted packets to pass. This pass-through arrangement means that the intermediate firewall is simply used as a router for this type of traffic.

Low-Medium

For all Virtual Private Connections over the Internet, COMPANY firewalls must operate in the Trusted Link mode, encrypting VPN traffic but not requiring the use of firewall proxies for VPN traffic.

High

For all Virtual Private Connections over the Internet, COMPANY firewalls must operate in the Private Link mode, encrypting VPN traffic and requiring the use of firewall proxies limiting the services available to remote VPN hosts.

5.4.2 Remote System Access

While Internet-based attacks get most of the media attention, most computer system break-ins occur via dial-up modems. As discussed in section 4.2.1, there are a variety of

configurations for supporting remote access via dial-up lines and other means. In general, the major security issue is authentication - making sure that only legitimate users can remotely access your system. The use of one-time passwords and hardware tokens is recommended for most companies.

Another issue is the organization's ability to monitor the use of remote access capabilities. The most effective approach is to centralize the modems into remote access servers or modem pools.

Low

All users who access the company system through dial-in connections periodically change their passwords. Direct dial-in connections to company production systems must be approved by the Network Services Manager.

Medium-high

All users who access the company system through dial-in connections must use one-time passwords. Direct dial-in connections to company production systems must be approved by the Network Services Manager and the Information Security Manager. The use of desktop modems to support dial-in access to company systems is prohibited.

Low-Medium-High

Information regarding access to company computer and communication systems, such as dial-up modem phone numbers, is considered confidential. This information must not be posted on electronic bulletin boards, listed in telephone directories, placed on business cards, or made available to third parties without the written permission of the Network Services Manager. The Network Services Manager will periodically scan direct dial-in lines to monitor compliance with policies and may periodically change the telephone numbers to make it more difficult for unauthorized parties to locate company communications numbers.

5.4.3 Access to Internal Databases

Another key architectural decision is how to securely support external access to internal databases. For small, relatively static user populations (i.e., company employees) that answer may be to create a VPN whose perimeter includes all users who need to access the internal database. For large or frequently changing user populations, such as the general public, some other means of securing access needs to be provided.

One solution is to host the database outside the company firewall. This poses the lowest risk to internal systems but makes the database vulnerable to outside attack and may require connectivity through the firewall to allow internal access or updates. This architecture is not workable in many cases, where the database needs to be accessed frequently by internal legacy systems.

Another approach is to provide connectivity from an external (outside the firewall) server through the firewall to an internal database. Many firewalls now support SQL proxies to limit the risk of providing this connectivity. The use of these proxies is fairly complex and requires training and expertise to assure a secure configuration. Allowing such

connections through the firewall greatly increases risk and should be avoided whenever possible.

Low-medium

All connections between external users and internal databases must be through the appropriate proxies on the COMPANY firewall. Only proxies approved by the Network Services Manager will be installed on the firewall.

High

No connections between external users and internal databases will be allowed. Where applications require Internet or Web access to COMPANY databases, the databases will be hosted external to the COMPANY firewall. Updates to these databases will be via "air gap" means, such as using floppy disks or other portable media. Any sensitive or confidential information to be stored on an external server will be encrypted.

5.4.4 Use of Multiple Firewalls

Virtual Private Networks is one example of the use of multiple firewalls. Bandwidth, availability or other performance requirements often dictate the use of multiple parallel firewalls, as well. Since the firewalls are in parallel, the security policy implemented by each firewall must be identical, or the resultant level of security provided would be that of the least secure firewall.

When multiple firewalls are used in parallel for availability or performance reasons, the configuration of each firewall shall be identical and under the control of a single firewall administrator. The Network Services Manager must approve any change to any firewall in this type of configuration.

Multiple internal firewalls may also be used to segment networks to provide access control to sensitive data and support detailed logging of such access. Known as "Intranet" firewalls, this type of configuration is often used by medium-high security organizations to wall off human resource or financial systems on internal networks.

Medium-high

Any sensitive or confidential information made accessible on internal networks that are connected to the Internet must be protected by a company-approved firewall. The firewall should be configured to limit access to such data to only authorized internal users.

Load balancing is the process whereby network traffic flowing through a network goes through more than one firewall in order to have a high network throughput. This is analogous to having more than one route to a city. The reason for doing this is twofold: load sharing and firewall backup (in case of the failure of one firewall, security will be maintained). Therefore, a site that wishes to include this policy might write one as follows:

More than one firewall may be deployed so that in case of one firewall failure, access into our network is still controlled by a firewall or so that network load is shared among these firewalls. In such a configuration the operational parameters of all firewalls shall be identical to assure a consistent security configuration.

5.5 Incident Handling

An incident is defined as an event that has actual or potentially adverse effects on computer or network operations resulting in fraud, waste, or abuse; compromise of information; or loss or damage of property or information. Examples include penetration of a computer system, exploitation of technical vulnerabilities, or introduction of computer viruses or other forms of malicious software (as defined by FIRST - Forum of Incident Response and Security Team).

Although the focus on securing a connection to the Internet is often on protecting from external threats, the misuse of the Internet connection by internal users is often a significant threat as well. Internal users may now have wide open access to internal databases via Virtual Private Networks or intranets that never existed before. Internal users may also be tempted to explore other systems over the Internet, causing your computer system to be the launching point for Internet attacks. Incident handling needs to address internal incidents as well as those instigated by external threats.

5.5.1 Intrusion Detection Overview

Intrusion detection plays an important role in implementing an organizational security policy. As information systems grow in complexity, effective security systems must evolve. With the proliferation of the number of vulnerability points introduced by the use of distributed systems, it is not enough to just "bolt the doors" and hope the locks hold. Some type of assurance is needed that the network is secure — that "all the doors are closed, they are strong and the locks are working." Intrusion detection systems can provide part of that assurance.

Intrusion detection provides two important functions in protecting information system assets. The first function is that of a feedback mechanism which informs the security staff as to the effectiveness of other components of the security system. In this sense intrusion detection is like a report card for perimeter defense sub-systems such as firewalls and dial-up access control systems. The lack of detected intrusions is an indication that the perimeter defenses are working, if a robust and effective intrusion detection system is in place. The second function is to provide a trigger or gating mechanism that determines when to activate planned responses to an incident.

A Framework for Action

"As the technology and policy to thwart attacks improves, however, so too does the ability of the attackers … new techniques in intrusion detection and incident response are needed. Centers to consolidate incident data and chart trends were invaluable in identifying the widely publicized "sniffer" attacks in 1994."

FEDERAL INTERNET SECURITY

The Federal Networking Council

This section presents an overview of different methods that may be used to detect intrusions into a computer information system. This is a representative list of the tools that various organizations are using today, hence there will always be new tools evolving. Not all of these tools are used within all network environments, nor should they be. Instead, the tools which are appropriate within the context of asset valuation, risk assessment, cost justification, and resources available should be selected for each situation.

Security is normally enforced through a combination of technical and traditional management methods. A decision must be made on the role technology will play in enforcing or supporting the policy. The methods listed tend to be technology based, although some discussion is provided of management tasks relative to these technologies.

The last portion of this section on intrusion detection contains some sample policy statements, ranging from lowest security to highest security afforded.

5.5.2 Methods

Intrusion detection can be implemented in different ways, and the choice of implementation method should be made based on the type of network being protected, the perimeter defense systems used, and the level of protection required by local policy. There are a number of methods for performing intrusion detection.

Regardless of the method chosen, an organization should have a defined Incident Response Capability. This may simply be a designated point of contact for users to report suspected incidents to, or it may be as formal as a team that uses proactive methods and tools to prevent incidents. The NIST Computer Security Handbook provides details on establishing and IRC, and the CERT and CIAC centers detailed in the resource system provide on-line capabilities. NIST Special Publication 800-3 also provides information on setting up an Incident Response Capability.

One method is *passively waiting* for complaints from users or others. Typical complaints might be that files have been modified or removed, or that disk volumes on servers are full for no apparent reason. The benefit of this method is that it is easy to implement. There are several disadvantages to this method. Obviously this method adds little protection to information systems or assurance to the security policy. Sophisticated attackers generally will not create such obvious symptoms. By the time it is apparent that the network has been attacked, it is too late to prevent damage. In extreme cases, the first indication that something is wrong may be when law enforcement investigators, or worse, newspaper reporters arrive on the premises.

Another method is *reviewing audit logs* periodically, searching for exceptional events. The events, which are of interest, could be excessive failed authentication attempts, excessive permission violations, unusual traffic patterns, etc. This method offers some additional protection beyond passive complaint based detection. Depending upon the frequency of the audit review, it may provide sufficient information to limit the impact of an attack. Generally, all modern computer and network systems will provide the required level of audit capability as a standard feature. Often, this feature is disabled by default and must be explicitly enabled. This method will require administrative efforts on an ongoing basis. Its effectiveness is dependent upon consistent and frequent manual review of log records. If the log function built into the underlying operating system or application is used, and that system or application is not sufficiently hardened against attacks, this method may be circumvented by sophisticated attackers who cover their tracks by disabling logging functions during their unauthorized activities, or by editing the log files to remove incriminating evidence.

Other *monitoring tools* can easily be constructed using standard operating system software, by using several, often unrelated, programs together. For example, checklists of file ownership's and permission settings can be generated and stored off-line. These

lists can then be reconstructed periodically and compared against the master. Differences may indicate that unauthorized modifications have been made to the system.

It is important to vary the monitoring schedule. System administrators can execute many of the commands used for monitoring periodically throughout the day. In addition, by running various monitoring commands at different times throughout the day, it becomes harder for an intruder to predict your actions. For example, if an intruder knows that each day at 5:00 p.m. the system is checked to see that everyone has logged off, he will simply wait until after the check has completed before logging in. But if the system administrator performs routine monitoring at random times of the day, an intruder cannot guess when this will happen, and thus a greater risk of detection.

Integrity-monitor tools for UNIX systems such as Tripwire perform a checksum of a file or file system, so that subsequent checksum results of the same system can be compared in order to detect modifications. These tools require installation by skilled system administrators. Ongoing administrative time is required to ensure that the integrity checks are consistent. Since the security mechanism is not part of the underlying operating systems or applications, it is much less likely that an attacker can cover the tracks left by their activities. Unfortunately, these tools may only be effective at indicating attacks where system modules are modified and may not detect other types of attacks, such as those where information is stolen by copying files.

Alarms and alerts from perimeter access control systems can indicate that a suspected attack is underway. Some perimeter access control systems such as firewalls or dial-up access control systems may be configurable to provide alarms if certain access rules are violated, error thresholds are exceeded, etc. These alarms may be audible, visual, email, pager, or messages to higher level management systems, i.e. SNMP traps. Once implemented, this type of detection may be relatively inexpensive from a management perspective, as the system may configured to send alerts to network management staff who are already monitoring other aspects of network status, i.e. dedicated personnel are not required. However, only intrusions that traverse the known perimeter systems will be detected. Intrusion from external networks via covert or unknown channels will not be detected, nor will unauthorized access of sensitive hosts or servers by employees or other valid network users. Another factor to consider is that if an attacker is capable of penetrating these perimeter access control systems, there is no assurance that they will not also disable any alarm functions provided by those systems.

Automated tools exist which perform *real-time analysis of data traffic*, and employ advanced logic to detect patterns of activity that indicate that an intrusion attack is underway. These tools can be host based, installed on each host system that is deemed critical, or network based, installed at centralized data traffic locations to allow all traffic to be monitored. These tools may be deployed so that attacks from both internal and external sources may be detected. Since they are independent of both host/server operating systems and perimeter access control systems, they are less likely to be subverted by attackers who successfully penetrate those systems. The success of these systems is dependent upon accurate foreknowledge of behavior patterns that indicate an intrusion, and this may not be possible. If overly specific patterns are defined, actual observed intrusive behavior may not match the target behavior. If less specific patterns are defined, excessive false alarms may result. This type of approach requires sophisticated heuristics that may overly complicate the use of the tool.

Tools also exist which regard statistical anomalies as a possible indication of an intrusion. This is done by maintaining a statistical profile of various network entities such as individual users, groups of users, applications, servers, etc., and then comparing the observed behavior of a member of one of these classes of entities. If the observed behavior falls outside of the range of the statistical profile, then this is an indication of a possible intrusion. This type of approach requires sophisticated heuristics that may overly complicate the use of the tool.

Use of sophisticated *software forensics* may identify authorship of various code modules. By routinely analyzing modules on protected systems, substitution of valid software by intruders can be detected. This esoteric approach theoretically offers protection against attacks which would not be detected by network perimeter defenses, such as those which use covert channels, or attacks by internal users where those users are knowledgeable and sophisticated enough to circumvent normal host security. The possible benefit of this method must be balanced against the normally low probability of such an attack and the complexity of the defense, as well as its limitation to detecting software modification such as introduction of Trojan horse programs.

5.5.3 Incident Response

The computer security policy should define the approach to be taken when dealing with a suspected intrusion. The procedures for dealing with these types of problems must be written down. A number of questions must be addressed before an incident occurs, so that the answers result from a calm, business-like consideration rather than the possible panic that may arise during the excitement of the incident. Some of the questions that must be addressed include:

- Who has authority to decide what actions will be taken?
- When should law enforcement be involved?
- Should your organization cooperate with other sites in trying to track down an intruder?
- Should the intrusion be stopped immediately upon detection, or should the suspected intruder be allowed to continue. By allowing the suspected intrusion to continue, additional evidence may be gathered in order to understand the method of attack to prevent a recurrence, as well as for possible use in tracking down the intruder(s) and for bringing civil and/or legal action against them.

Answers to these questions should be part of the incident handling procedures. If incident handling procedures are not in place, they need to be developed. The intrusion detection systems and procedures outlined here are only one part of an overall security program. While some limited utility may be derived from any one component of a security program (access control, intrusion detection, incident response, etc.) for best results, all components should be implemented in a unified approach based on a security policy developed for the specific site. For example, if server-based alarms are forwarded to the client/server support group, but firewall alarms are handled by the network support group, the extent of an intrusion may be underestimated or missed entirely.

Intrusion Detection Policy - Low Risk

Implementation:

Operating system and application software logging processes shall be enabled on all host and server systems.

Alarm and alert functions, as well as logging, of any firewalls and other network perimeter access control systems shall be enabled.

Administration:

System integrity checks of the firewalls and other network perimeter access control systems must be performed on a routine basis.

Audit logs from the perimeter access control systems shall be reviewed daily.

Audit logs for servers and hosts on the internal, protected network shall be reviewed on a weekly basis.

Users shall be trained to report any anomalies in system performance to their system administration staff, as well as relevant network or information systems security staff.

All trouble reports received by system administration personnel should be reviewed for symptoms that might indicate intrusive activity. Suspicious symptoms should be reported to Network or Information Systems security personnel.

Intrusion Detection Policy - Medium Risk

Implementation:

Normal logging processes shall be enabled on all host and server systems.

Alarm and alert functions, as well as logging, of any firewalls and other network perimeter access control systems shall be enabled.

All critical servers shall have additional monitoring tools such as tripwire or appropriate software wrappers installed, as a supplement to the activity logging process provided by the operating system. Examples: Domain Name Servers, authentication servers, security servers in the Unix environment, domain controllers and Exchange servers in the Windows NT environment, and any application server which is considered to be mission critical should be afforded this protection.

Administration:

System integrity checks of the firewalls and other network perimeter access control systems must be performed on a routine basis.

Audit logs from the perimeter access control systems shall be reviewed daily.

Audit logs for servers and hosts on the internal, protected network shall be reviewed on a weekly basis.

User education shall be provided in order to train end users of computing systems to report any anomalies in system performance to their system administration staff, as well as relevant network or information systems security staff.

All trouble reports received by system administration personnel should be reviewed for symptoms that might indicate intrusive activity. Suspicious symptoms should be reported to Network or Information Systems security personnel.

Host based intrusion tools such as tripwire will be checked on a routine basis.

The IRC or network security personnel will establish relationships with other incident response organizations, such as other IRCs within the organization or FIRST (FIRST maintains contact with many IRCs – see www.first.org.) and share relevant threats, vulnerabilities, or incidents.

Unless critical systems have been compromised, the organization will first make an attempt to track intruders before correcting systems. (Name of specific person or office) has the authority to make decisions concerning closing security holes or attempting to learn more about the intruder. This person must be well training in legal issues surrounding incident handling.

Intrusion Detection Policy - High Risk

Implementation:

Normal logging processes shall be enabled on all host and server systems.

Alarm and alert functions, as well as logging, of any firewalls and other network perimeter access control systems shall be enabled.

All servers shall have additional monitoring tools such as tripwire or appropriate software wrappers installed, as a supplement to the activity logging process provided by the operating system.

All critical servers shall have redundant intrusion detection tools installed, which operate on a different principle from the primary tool that is installed on all servers. For examples: If the primary IDS tool is tripwire, which uses a checksum based approach to ensuring system integrity, then the critical servers will also have expert systems IDS tools which use a statistical anomaly approach.

At logical network concentration points, IDS tools will be installed which monitor for traffic patterns consistent with known attacks.

Administration:

System integrity checks of the firewalls and other network perimeter access control systems must be performed on a routine basis.

Audit logs from the perimeter access control systems shall be reviewed daily.

Audit logs for servers and hosts on the internal, protected network shall be reviewed on a daily basis.

User education shall be provided in order to train end users of computing systems to report any anomalies in system performance to their system administration staff, as well as relevant network or information systems security staff.

All trouble reports received by system administration personnel should be reviewed for symptoms that might indicate intrusive activity. Suspicious symptoms should be reported to Network or Information Systems security personnel.

Host based intrusion tools such as tripwire, will be checked on a daily basis.

Network traffic monitoring IDS systems will be checked on a periodic basis for proper function and configuration.

The IRC or network security personnel will establish relationships with other incident response organizations, such as other IRCs within the organization or FIRST (FIRST maintains contact with many IRCs – see www.first.org.) and share relevant threats, vulnerabilities, or incidents.

The organization will attempt to prosecute intruders but will not allow security holes to go uncorrected in order to learn more about the intruder.

5.6 Administrative

Internet security policy must be closely integrated with the day-to-day use of the Internet by its users and the day-to-day management of the network and computer systems. It must also be adopted into the organization's culture and environment through education (see Section 5.7). This document concentrates on areas of technical policy. However, administrative decisions such as assigning security responsibility are, often, more important in the long run. This chapter addresses additional aspects of Internet security that need to be considered as administrative issues.

In addition to the myriad technical and administrative responsibilities of the network administrator for internal operations, there are responsibilities related to Internet connectivity. This chapter covers areas not addressed in other chapters.

- Assigning security responsibility.
- Resolving violations/establishing penalties.
- Appropriate use, including restricting access to specific WWW sites or Usenet newsgroups, etc., in accordance with company policy. (Appropriate use of email is covered in the chapter on Email.)
- Establishing a privacy policy, specifically including email privacy and network monitoring. (Email privacy is including in the chapter on Email.)

5.6.1 Assigning Security Responsibility

The success of any security policy depends more on the motivation and skill of the people administering the policy than it does on any sophisticated technical controls. Effective Internet security begins with the network administrator(s) (often called the LAN or System administrator).

The network administrators are responsible for implementing LAN security as it pertains to the Internet. If there are multiple network administrators, it is important that roles be coordinated. For example, an attack on a company's Web site may involve contingency plans for response and restoring to be carried out by a Web administrator as well as an increase in the local auditing and monitoring of the internal network by the LAN administrator and increased surveillance of the firewall by a Firewall administrator. It is an organizational decision, often having to do with size, whether or not to group various network administration functions together.

The organization should specifically name the person or office responsible for the day to day security of the Internet connection. This duty will often be given to the

network or LAN administrator but may also be given to a separate security organization. In this case, it is imperative that the network administrator and the security officer coordinate closely and that the security officer be well-versed in Internet protocols.

The person or office responsible for Internet security is _____.

The person responsible for Internet security may have considerable technical power to configure the firewall, create and suspend userIDs and review audit reports. The actions of this person should be monitored by using separation of duties (if there are multiple network administrators) or through careful screening of the individual. A disgruntled or malicious network administrator is a significant problem.

Higher risk organizations may wish to use the following policies.

A personnel screening process should be in place for critical roles. Critical roles such as system and LAN administrators, security personnel, and other sensitive positions as determined by senior management must satisfactorily complete the screening before being given system manager privileges.

System and LAN administrators and other privileged roles are given incremental access. In other words, newly hired system administrators are not given full system privileges if their job function does not require it. More privileges are given as the breath of their job function increases, and consequently, their level of trust.

Managers of critical roles are responsible for determining job function and scope so that broad system and network privileges are not excessively given.

Unfortunately, system policy will not always be followed. For any of the specific policies discussed in this guide it may be appropriate for an organization to state penalties for non-compliance. For the most part, this is only necessary if the penalty is severe or could be construed as severe. (For example, many organizations have stated penalties which are severe for software copyright violations because employees viewed the act as "minor" whereas as the organization and the copyright holder viewed it as "major.") Other violations can be dealt with on a case-by-case basis using the procedures the organization uses for other personnel problems.

The LAN Administrator may temporarily suspend access privileges of any user if deemed necessary to maintain the integrity of the computer or network. (A more strict policy may cite the need for authorization from the Security Officer or System Manager prior to the system administrator taking or restoring system privileges.)

5.6.2 Appropriate Use

Similar to policies for appropriate use of the telephone, organizations need to define appropriate use of the Internet and the World Wide Web. While it is tempting to simply state that any such use must be for business purposes only, it is generally recognized that this type of policy is completely unenforceable. If a policy cannot be consistently enforced, non-compliance is inevitable and the policy will have no force as a basis for punitive action. Acceptable use of electronic mail as well as the related privacy concerns is identified in 6.3.1.

Higher risk organizations that cannot tolerate a more flexible attitude toward casual use of the Internet might consider some alternative solutions that may fit better into the company culture and are enforceable:

A separate public access server with a commercial Internet service provider or other link to the Internet could be set up for employee use. This service would not be connected to any internal systems and could be used for incidental purposes in accordance with the company's appropriate use policy.

Software tools such as a firewall can be used to block access to all Internet sites except those that the organization has approved.

No matter what the risk environment, certain uses of the company connection to the Internet can never be sanctioned. The use of the company connection to the Internet is inappropriate when that use:

- Compromises the privacy of users and their personal data.
- Damages the integrity of a computer system, or the data or programs stored on a computer system.
- Disrupts the intended use of system or network resources.
- Wastes resources that are needed for business use (people, network bandwidth, or CPU cycles).
- Uses or copies proprietary software when not authorized to do so.
- Uses a computer system as a conduit for unauthorized access attempts on other computer systems.
- Uses a government, corporation, or university-owned system for private purposes or for purposes not in the direct interest of the government, corporation, or university.
- Consists of unauthorized and excessive snooping, probing, or otherwise connecting to a node or nodes in a manner that is deemed not to be of an authorized nature.
- Results in the uploading, downloading, modification, or removal of files on any node in the network for which such action is not authorized.

Whether an organization's policy on acceptable use of the Internet is flexible, tolerant, or severe, education is the key to influencing user behavior. Users need to be educated on how they are a part of the company's reputation and how Internet use can affect that reputation. Users should know that every visit to an Internet site leaves a "footprint" and why the company reserves the right to monitor its resources. Administrators need to know their responsibilities regarding the implementation of tools and technical controls (e.g., blocking access to sites, monitoring, and reporting) - so that the company's acceptable use policy can be enforced.

Internet Usage Policy - Low Risk

The Internet is considered a valuable company asset. Users are encouraged to make use of the Internet and explore its uses. With such open access, employees must maintain a diligent and professional working environment.

Employees may not use the Internet for personal commercial purposes, may not access any obscene or pornographic sites, and may not access or use information that would be considered harassing. Employees abusing such privileges will be subject to monitoring of their computer system activity and disciplinary action, ranging from verbal reprimands to termination or legal prosecution.

Access to the Internet from a company-owned home computer or through company-owned connections must adhere to all the same policies that apply to use from within company facilities. Employees should not allow family members or other non-employees to access company computer systems.

Users posting to Usenet newsgroups, Internet mailing lists, etc. must include a company disclaimer as part of each message.

It is impossible to define all possible unauthorized use, therefore disciplinary action may occur after other actions if the circumstances warrant it. Examples of other behavior deemed unacceptable which would result in disciplinary action include:

Unauthorized attempts to break into any computer.

Using company time and resources for personal gain.

Theft or copying electronic files without permission.

Sending or posting company confidential files outside the company or inside the company to unauthorized personnel.

Refusing to cooperate with a reasonable security investigation.

Sending chain letters through e-mail.

Internet Usage Policy - Medium Risk

A company's computer systems and networks are provided for business use only. Occasional, reasonable personal use is allowed. Any use perceived to be illegal, harassing, offensive, in violation of other company policies, or any other uses that would reflect adversely on the company can be the basis for disciplinary action up to and including termination or judicial action. All employees are expected to conduct their use of these systems with the same integrity as in face-to-face or telephonic business operations.

Another approach to stating the above might be:

Company communications systems and equipment, including electronic mail and Internet systems, along with their associated hardware and software, are for official and authorized purposes only. Managers may authorize incidental use which: does not interfere with the performance or professional duties; is of reasonable duration and frequency, serves a legitimate company interest, such as enhancing professional interests or education, and does not overburden the system or create any additional expense to the company.

Users posting to Usenet newsgroups, Internet mailing lists, etc. must include a company disclaimer as part of each message.

Personal accounts on on-line services should not be used from company computers. A company subscription for commercial Internet services or fee-for-use services must be in place prior to using company-owned equipment to access these commercial services.

Passwords to company systems are provided in order to protect sensitive information and messages from unauthorized use or viewing. Such passwords are not intended to prevent appropriate review by company management. Company management reserves the right to periodically monitor employees' use of any computer systems or network.

Managers are responsible for ensuring that assigned personnel understand Internet acceptable use policy.

Access to the Internet from a home computer must adhere to all the same policies that apply to use from within company facilities. Employees should not allow family members or other non-employees to access company computer systems.

Internet Usage Policy - High Risk

COMPANY is fully interconnected with the Internet and other networks. In general, valid users enjoy unrestricted network access. However, access from the Internet or other sites to or through company Internet resources is only authorized when that access is in conjunction with valid work or project-related requirements.

A separate public access server to the Internet is provided for employee's personal use. This service is to be used with professional discretion. Only those sites to which the company has approved access are available through this service.

All employees are expected to conduct their use of these systems with the same integrity as in face-to-face or telephonic business operations. Any use perceived to be illegal, harassing, and offensive or in violation of other company policies can be the basis for disciplinary action up to and including termination or judicial action.

Users posting to Usenet newsgroups, Internet maillists, etc. must include a company disclaimer to their message.

Company management reserves the right to periodically monitor employees' use of any computer systems or network.

Access to the Internet from a home computer must adhere to all the same policies that apply to use from within company facilities. Employees should not allow family members or other non-employees to access company computer systems.

5.6.3 Privacy

The privacy policy for Internet usage should be consistent with other privacy policies. Just because it is (technically) easy to monitor employees does not mean it is always a good idea. Employee morale is a significant factor to security, as well as productivity. Employees should be made aware that network records are subject to release for conditions outside the organization's control, such as a subpoena or, for Government agencies, a Freedom of Information Act request.

Low and medium risk

The Internet connection is an organization resource. Activities may be subject to monitoring, recording, and periodic audits to insure they are functioning properly and to protect against unauthorized use. In addition, the organization may access any user's computer accounts or communication. The organization will disclose information obtained through such auditing to appropriate third parties, including law enforcement authorities or FOIA requesters. Use of (resources) is expressed consent by the user to such monitoring, recording and auditing.

5.7 Awareness and Education

Most companies' computer users generally fall into one of three camps: those that are Internet "wizards"; those that are somewhat knowledgeable, but haven't had much experience with it; and those that have heard of it, know that great stuff is out there, but have no notion as to how to proceed.

Most users are generally aware that there are security risks related to Internet use, but don't necessarily understand what the security issues are. They often do not know how to recognize a security problem or how to include Internet security procedures (rules of behavior) into their daily computer lifestyles. They may not know the consequences for inappropriate or unauthorized actions.

Making computer system users aware of their security responsibility and teaching them correct practices helps users change their behavior. Users cannot follow policies they do not know about or understand. Training also supports individual accountability, which is one of the most import ways to improve computer security. Without knowing the necessary security measures and how to use them, users cannot be truly accountable for their actions. Training is also needed for network administrators who need special skills to understand and implement the technology needed to secure Internet connections. The risks of Internet connectivity should be emphasized to upper management to ensure their support is obtained.

All users, managers, and administrators who are given Internet access should receive initial and periodic security awareness and training appropriate for their use of the Internet. Training for experienced users should focus on acceptable use issues. The reasons why the company has adopted a policy, for instance, against certain rec newsgroups needs to be understood. Personal postings or replies to postings that include the "company.com" header, whether disclaimers are used or not, immediately reflect the company. This might be a shock to such users if they have come from a university setting. Emphasis should be placed on roles and responsibilities, along with technical security issues. Technical administrators should be trained in their responsibilities for implementing technical policy on their system and networks as described in section 5.6.1.

Users that are somewhat knowledgeable need to be educated to the methods and ways of the Internet. These users may be more familiar with bulletin board type systems and must be informed that their postings now span worldwide, rather than the US or a local community. These users need to be patient: they should browse and watch for awhile before jumping in to a discussion group. They should not download software or subscribe to information until the knowledge base improves—it could cause the company to be at risk. Somewhat knowledgeable users need to be educated as to how to become experienced Internet users in addition to the acceptable use, responsibilities, technical and security issues.

Inexperienced users need the full array of help. They must be educated as to what the Internet is, kinds of services are available, what the community is, and how to interact with it, They must learn about Internet mailing lists, newsgroups, search engines and etiquette on the Internet to name a few. They must also receive the education given to experienced and partly knowledgeable users.

Internet Security Education - Low Risk

It is the policy of this organization to provide periodic security awareness training to all managers, operators, and end users as described in the education section of the Company's Policy document. Such training will be augmented with the very new and rapidly changing issues regarding security and the Internet. Users are encouraged to

scan security-related lists, keep up with security issues and technology and share security relevant information with the Security Dept.

Internet Security Education - Medium Risk

There is a fairly liberal culture on the use of the Internet; however, more knowledge needs to be gained on how to effectively use the Internet capabilities and its security implications. The following policies are in addition to the policy listed for “low risk” organizations.

Internet security training shall include combined or separate training sessions from the regular security education curriculum, news flashes or tips via the system, memos, computer incident alerts, and other appropriate training as determined by the organization. Issues and topics covered may include firewall training, downloading of information and software, Applet software (Java, ActiveX), email, mail lists, home pages, browsers, acceptable use of Internet, etc.

Network and firewall administrators and staff, and technical managers of networks with an Internet connection shall receive training on managing network security, tradeoffs, and costs to various approaches, kinds of attacks that can occur, network architecture, and security policy issues.

LAN and system administrators shall receive technical hands-on Firewall training,

Certain system and network configuration and scanning tools (e.g., pingall, SATAN) have become a must for security reviews in terms of identifying active systems and IP address, current configuration parameters, etc. Additionally virus and vulnerability scanning tools both in the public and commercial domain (e.g., SATAN, ISS, NETProbe, PINGWARE, COPS, Tripwire, etc.) are extremely useful for exposing system vulnerabilities. All network and system administrators will be educated in their uses. They will also be required to keep up-to-date on such technologies.

New users will receive an orientation to the Internet which will include hands-on training, and and a review of its security considerations. All users will sign an Internet Acceptable Use agreement.

Internet Security Education - High Risk

The organization would like to increase the availability of the Internet to its users but in a conservative manner and only after the users are knowledgeable and competent regarding Internet security issues. The following policies are in addition to the policy listed for “medium risk” organizations.

Users will also receive continuous security training in the form of news flashes, security alerts or tips via the system, memos, computer incident alerts, and other appropriate training as determined by the organization. Issues and topics covered may include firewall training, downloading of information and software, Applet software issues (Java, ActiveX), e-mail, mail lists, home pages, browsers, acceptable use of Internet, encryption, etc.